

06-21-00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: Scott P. Teanican and Bryan D. Skene
Docket: 50002.4USU1
Title: METHOD AND SYSTEM FOR NETWORK LOAD BALANCING WITH A COMPOUND DATA STRUCTURE

JC490 U.S. PTO
09/596796 A
06/19/00

CERTIFICATE UNDER 37 CFR 1.10

'Express Mail' mailing label number: EL626731366US

Date of Deposit: June 19, 2000

I hereby certify that this paper or fee is being deposited with the United States Postal Service 'Express Mail Post Office To Addressee' service under 37 CFR 1.10 and is addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231.

By:

Name: Janice Kniep

BOX PATENT APPLICATION
Assistant Commissioner for Patents
Washington, D.C. 20231

Sir:

We are transmitting herewith the attached:

- ☒ Transmittal sheet, in duplicate, containing Certificate under 37 CFR 1.10.
- ☒ Utility Patent Application: Spec. 23 pgs; 21 claims; Abstract 1 pg.:
- ☒ 10 sheets of formal drawings with 12 FIGURES
- ☒ A copy of the verified statement to establish small entity status filed in the related provisional application
- ☒ An unsigned Combined Declaration and Power of Attorney
- ☒ Return postcard
- ☒ **PAYMENT OF THE FILING FEE IS BEING DEFERRED.**

MERCHANT & GOULD P.C.

P.O. Box 2903, Minneapolis, MN 55402-0903

(612) 332-5300

By:

Name: John W. Branch

Reg. No.: 41,633

Initials: JWB/jk



23552

PATENT TRADEMARK OFFICE

Applicants: Scott P. Tennican et al.

Attorney Docket No: FFIV114093

Serial No: 60/140,272

Filed: 6/18/99

Title: METHOD AND SYSTEM FOR NETWORK LOAD BALANCING WITH A
COMPOUND DATA STRUCTURE

VERIFIED STATEMENT (DECLARATION) CLAIMING SMALL ENTITY STATUS
(37 C.F.R. §§ 1.9(f) and 1.27(c)) - SMALL BUSINESS CONCERN

I hereby declare that I am

- () the owner of the small business concern identified below:
(X) an official of the small business concern empowered to act on behalf of the concern
identified below:

NAME OF CONCERN F5 Networks, Inc.

ADDRESS OF CONCERN 200 First Avenue West, Suite 500, Seattle, WA
98119

I hereby declare that the above-identified small business concern qualifies as a small business concern as defined in 37 C.F.R. §§ 121.3-18, and reproduced in 37 C.F.R. §§ 1.9(d), for purposes of paying reduced fees under 35 U.S.C. §§ 41(a) and (b) in that the number of employees of the concern, including those of its affiliates, does not exceed 500 persons. For purposes of this statement, (1) the number of employees of the business concern is the average over the previous fiscal year of the concern of the persons employed on a full-time, part-time or temporary basis during each of the pay periods of the fiscal year, and (2) concerns are affiliates of each other when either, directly or indirectly, one concern controls or has the power to control the other, or a third party or parties controls or has the power to control both.

I hereby declare that rights under contract or law have been conveyed to and remain with the small business concern identified above with regard to the invention, entitled METHOD AND SYSTEM FOR NETWORK LOAD BALANCING WITH A COMPOUND DATA STRUCTURE by inventors Scott P. Tennican and Bryan D. Skene, described in

- (X) Provisional Application Serial No. 60/140,272, filed 6/18/99.

If the rights held by the above-identified small business concern are not exclusive, each individual, concern or organization having rights to the invention is listed below* and no rights to the invention are held by any person, other than the inventor, who could not qualify

LAW OFFICES OF
CHRISTENSEN O'CONNOR JOHNSON & KINDNESS^{LLC}
1420 Fifth Avenue
Suite 2800
Seattle, Washington 98101
(206) 682-8100

as a small business concern under 37 C.F.R. § 1.9(c) or by any concern which would not qualify as a small business concern under 37 C.F.R. § 1.9(d) or a nonprofit organization under 37 C.F.R. § 1.9(e). *NOTE: Separate verified statements are required from each named person, concern or organization having rights to the invention averring to their status as small entities. (37 C.F.R. § 1.27)

NAME: NONE

ADDRESS: _____

- ☐ INDIVIDUAL
☐ SMALL BUSINESS CONCERN
☐ NONPROFIT ORGANIZATION

I acknowledge the duty to file, in this application or patent, notification of any change in status resulting in loss of entitlement to small entity status prior to paying, or at the time of paying, the earliest of the issue fee or any maintenance fee due after the date on which status as a small entity is no longer appropriate. (37 C.F.R. § 1.28(b))

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful, false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001, and that such willful, false statements may jeopardize the validity of the application, any patent issuing thereon, or any patent to which this verified statement is directed.

NAME OF PERSON SIGNING: Joann Reiter

TITLE OF PERSON OTHER THAN OWNER: General Counsel

ADDRESS OF PERSON SIGNING: 200 First Avenue West, Suite 500, Seattle, WA 98119

SIGNATURE:  DATE: 7-27-99

JWB:cah

LAW OFFICES OF
CHRISTENSEN O'CONNOR JOHNSON & KINDNESS^{LLC}
1420 Fifth Avenue
Suite 2800
Seattle, Washington 98101
(206) 682-8100

**METHOD AND SYSTEM FOR NETWORK LOAD BALANCING WITH A
COMPOUND DATA STRUCTURE**

Related Application

5 This application claims the benefit under 35 U.S.C. § 119(e) of U.S. Provisional
Application No. 60/140,272, filed June 18, 1999.

Field of the Invention

10 The present invention relates to methods of implementing different types of data
structures in database systems, and in particular, to a complex data structure for
retrieving data from a database.

Background of the Invention

15 Data such as letters, numbers, and words are often maintained and stored in
different types of data structures that complement their eventual use with an application
program. Since a List data structure enables the storing of data objects in an ordered
set, this data structure is typically employed with an application program that often
retrieves several sequentially related data objects at a time. For example, a multi-media
player program usually generates streaming video from sequentially ordered data
objects in a List data structure. Although retrieval times for large numbers of
sequentially ordered or related data objects can be improved by the use of a List data
20 structure, it may not always be the optimal data structure for retrieving individual and/or
several non-sequential or unrelated data objects from a large set of data objects. It is
well known that the maximum amount of time to locate an individual data object on a
list increases as the number of data objects referenced by the list increase.

25 Another type of data structure is the Trie, which is often used with an
application program that requests individual and/or several non-sequential or unrelated
data objects. A Trie stores data in each transition between each node in a multi-level

data structure, rather than at the node itself. In this way, all of the transitions in the path between the root and each leaf represent the data object (key) and each transition between each node in the path is associated with a single character/number of the key. Since the nodes themselves are unlabeled, each transition between each node is labeled with a particular character/number. Also, the Trie data structure facilitates the calculation of a constant maximum amount of time to retrieve a stored data object based on the number of levels (nodes associated with the number of alphanumeric characters in the longest data object) in the Trie, e.g., $O(1)$. Thus, the maximum search time for a given number of levels in a Trie data structure remains relatively constant as the actual number of data objects stored in the data structure increases.

For example, a dictionary program often employs a Trie data structure to store words and provide relatively constant maximum search times as the number of words (but not their length) in the dictionary increase. In a Trie data structure for a dictionary program, every character in the alphabet (A through Z) is partitioned into individual and disjoint main level search nodes. Also, the total number of search node levels in a Trie data structure for a dictionary corresponds to the number of characters in the dictionary's longest possible word.

Although the List and Trie data structures can complement the operation of particular application programs, there are many other types of application programs that can generate requests for both individual/unrelated and sequential/related data objects. A facility that could employ a particular request generated by an application program to choose a type of data structure most suited to reference and retrieve requested data would be an improvement over the prior art. Also, a data structure that could provide access by one or more unique keys to the same data and enable the automatic removal of all references to deleted data would complement the operation of application programs that generate different types of requests.

A more complete appreciation of the invention and its improvements can be obtained by reference to the accompanying drawings, which are briefly summarized below, to the following detail description of presently preferred embodiments of the invention, and to the appended claims.

Summary of the Invention

In accordance with the invention, the above and other problems are solved by employing a plurality of data structures to optimize the retrieval of at least one data object over a network. Each data object is stored in a data store and each data object is separately referenced in each of the plurality of data structures. In response to a request for one data object, one of the plurality of data structures is automatically determined to be best suited to retrieve the one data object. The determined data structure is employed to locate and retrieve the one data object from the data store. Also, in response to a request for a plurality of related data objects, another one of the plurality of data structures is automatically determined to be best suited to retrieve the plurality of related data objects. The other one of the plurality of data structures is employed to locate and retrieve the plurality of related data objects from the data store. Additionally, in response to a request to delete at least one data object, each reference to each deleted data object in the plurality of data structures is automatically deleted.

In accordance with other aspects of the invention, a parent object is associated with each data object. Each parent object identifies each reference for the associated data object in the plurality of data structures. When a data object is deleted, the parent object associated with the deleted data object is employed to identify each reference for the deleted data object in the plurality of data structures and each reference is deleted.

In accordance with still other aspects of the invention, the data object can be a collector object that is associated with a member object that identifies one or more other data objects that are related to the collector object. The member object is employed to reference and retrieve each data object related to the collector object when the collector object is retrieved.

In accordance with yet other aspects of the invention, the plurality of related data objects have at least one related characteristic, including port, IP address and type. Also, the plurality of data structures may include List, Hash and Trie.

In accordance with other aspects of the invention, one of the plurality of data structures can be a Trie data structure that identifies a key in the request for a data object. The Trie data structure divides the key into segments and each segment is

employed to search the Trie data structure and locate a reference to the requested data object. The key can represent an IP address and/or a port. Also, each segment can be represented by at least one bit.

In accordance with still other aspects of the invention, the data store is a
5 database. The database can be a relational, object-oriented, or some combination of the two, database. Also, the data store can be a data warehouse.

The invention may be implemented as a computer process, a computing system
or as an article of manufacture such as a computer program product or computer
readable medium. The computer program product may be a computer storage medium
10 readable by a computer system and encoding a computer program of instructions for
executing a computer process. The computer program product may also be a
propagated signal on a carrier readable by a computing system and encoding a computer
program of instructions for executing a computer process.

These and other features as well as advantages, which characterize the
15 invention, will be apparent from a reading of the following detailed description and a
review of the associated drawings.

Brief Description of the Drawings

FIGURE 1 illustrates a block diagram of a client-server operating
environment used in an embodiment of the invention;

20 FIGURE 1B illustrates a block diagram of a client-server operating
environment used with a server array controller in another embodiment of the invention;

FIGURE 2 shows a block diagram of a computing device used in an
embodiment of the invention;

FIGURE 3 illustrates a block diagram of an overview of the complex
25 data structure that includes a List data structure and a Trie data structure for one
embodiment of the invention;

FIGURE 4A shows a block diagram of a data object;

FIGURE 4B illustrates a block diagram of a collector object;

FIGURE 5 shows a block diagram of nested collector and data objects;

FIGURE 6 illustrates a block diagram of an overview of the complex data structure with greater detail shown for the List data structure referencing data objects and collector objects;

FIGURE 7 shows a block diagram of an overview of the complex data structure with greater detail shown for the Trie data structure referencing data objects;

FIGURE 8 illustrates a block diagram of an overview of the complex data structure with greater detail shown for the Trie data structure referencing data objects and collector objects;

FIGURE 9A shows a block diagram of an IP address list;

FIGURE 9B illustrates a block diagram of a Port list;

FIGURE 10 shows a block diagram of an eight level Trie data structure shown in greater detail;

FIGURE 11 is a flow diagram illustrating the operational characteristics for determining one of a plurality of data structures best suited to reference and retrieve a requested data object; and

FIGURE 12 is a flow diagram illustrating the operational characteristics for deleting each reference to a deleted data object in each of the plurality of data structures.

Detailed Description of the Preferred Embodiment

An embodiment of the invention provides for employing a complex data structure to optimize the retrieval of data from a data store over a network. The complex data structure may include multiple data structures, e.g., List and Trie, which in parallel separately reference the same data objects in a data store. For a particular functional request to retrieve data, the complex data structure examines the request to determine whether it is associated the List or Trie data structures. In most cases, the Trie data structure is associated with a functional request for a single data object and the List data structure is associated with functional request to retrieve several related data objects.

Each data object is associated with a parent object that includes a list of every reference to the data object in both the Trie and List data structures. When a data object

is subsequently deleted, the complex data structure employs the parent object list to automatically delete every reference to the deleted data object in both the Trie and List data structures. Additionally, the complex data structure provides for a particular type of data object, i.e., a collector object, that is associated with a member object and which includes a list of other related data/collector objects that are referenced in a sub-tree below the node level of the collector object in the Trie data structure. Also, when the data associated with a collector object is requested, other data associated with the other data/collector objects on the member object list are automatically retrieved.

Another embodiment of the invention may employ a complex data structure that includes a List data structure and Hash data structure, which can be used with the List data structure to process functional requests for single objects. Generally, the Hash data structure requires less overhead, e.g., memory and processor cycles, than a Trie data structure to process a request for data. However, the Hash data structure would not support collector objects as discussed in greater detail below.

The logical operations of the various embodiments of the invention are implemented (1) as a sequence of computer implemented actions or program modules running on a computing device; and/or (2) as interconnected hardware or logic modules within the computing device. The implementation is a matter of choice dependent on the performance requirements of the computing system implementing the invention.

Accordingly, the logical operations making up the embodiments of the invention described herein are referred to alternatively as operations, actions or modules. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Typically, the functionality of program modules may be combined or distributed in various embodiments.

FIGURE 1A illustrates one embodiment of an overview of a suitable client-server operating environment in which the invention may be implemented. A client 12 is in communication with a server 16 over a network 14 such as the Internet. The server 16 is in communication with a data store 18 for retrieving data that may be requested by the client 12. In response to a request from the client 12 over the network

14, the server 16 can access the data store 18 and provide the requested data in an appropriate format to the client.

FIGURE 1B illustrates another embodiment of an overview 40 of an exemplary operating environment in which the invention may be implemented. A network device 42 is in communication with a server array controller 46 over a network 44 such as the Internet. The network device 42 may be a server, client, router, firewall, cache or another server array controller and the network 44 may also be a local area network or a wide area network for an enterprise. The server array controller 46 is in communication with a data store 48 that stores the address for at least one of a pool of virtual servers 52 managed by the server array controller 46. The virtual servers 52 are in communication over another network 50 with the server array controller 46, which manages the requests for data from the virtual servers. The server array controller 46 uses the information in the data store 48 to transform and direct traffic from one network device to another. The other network 50 may be a local area network, wide area network or the Internet.

FIGURE 2 illustrates one embodiment of a suitable operating environment for a computer 20 in which the invention may be implemented. The computer 20 is only one embodiment of a suitable operating environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Other well known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, personal computers, server computers, client computers, laptop devices, multiprocessor devices, microprocessor based systems, programmable consumer electronics, network PCs, minicomputers, mainframe computers, client-server environments, distributed computing environments that include any of the above systems or devices, and the like.

FIGURE 2 shows functional components of the computer 20, including a processor system 28, a CPU 30, a memory system 22, a network interface 34, a computer readable media device 32 and an input/output device 36. The memory system 22 generally includes both volatile memory (e.g., RAM) and non-volatile memory (e.g., ROM, PC cards, etc.). An operating system 26 is resident in the memory system 22 and

executes on the CPU 30, such as Unix®, Linux®, or the Windows NT® operating system from the Microsoft ® Corporation.

One or more application programs 24 are loaded into the memory system 22 and run on the operating system 26. Examples of applications include email programs, scheduling programs, network management programs, PIM (personal information management) programs, word processing programs, spreadsheet programs, Internet browser programs, and so forth. The computer 20 has a power supply (not shown), which can be implemented as one or more batteries or might further include an external power source that overrides or recharges built-in batteries, such as an AC adapter.

The computer 20 may operate with at least some form of computer readable media. Computer readable media can be any available media that can be accessed by the computer readable media device 32. By way of example, and not limitation, computer readable media may comprise computer storage media and communications media. Computer storage media includes volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules, or other data. Computer storage media includes, but is not limited to, random access memory (RAM), read only memory (ROM), electrically erasable programmable read only memory (EEPROM), flash memory or other memory technology, CD-ROM, digital versatile disk (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by the computer 20. Communication media typically embodies computer readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term “modulated data signal” means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, radio

frequency (RF), infrared and other wireless media. Combinations of any of the above should also be included within the scope of computer readable media.

FIGURE 3 shows an overview 100 of a complex data structure employed by the invention to provide optimal retrieval of requested data from a data store 108. A
5 Functional Interface 102 processes each request for data. The functional interface is in communication with a List data structure 104 and a Trie data structure 106 and it determines which data structure matches the particular functional request for data. The Trie 104 and List 106 data structures provide parallel access to data stored in the data store 108. Although not shown, another embodiment of the invention may employ a
10 Hash function data structure (instead of a Trie data structure) with the List data structure to optimize the retrieval of data in response to a functional request.

Objects in the data store 108 are retrieved/added/deleted through the Functional Interface 102. Retrieval functions are specific to either the Trie data structure 104 or the List data structure 106. For the Trie data structure 104, there is a Trie retrieval
15 function that accepts a top level accessor object and a unique key and returns a container object, e.g., a data object or a collector object which are discussed in greater detail below. The List data structure 106 supports several retrieval functions, including mid level list accessor objects, changing the current state of a list, and retrieving at least one of the container objects that is a member of a list. Since the mid level list accessor
20 objects reference associated top level Trie accessor object, they do not need to be specified by list functions.

A Trie retrieval function is generally employed to find a single container object when at least one of the Trie keys that give access to the object are known. For example, a Trie retrieval occurs when a data packet arrives at an appliance on an
25 interface where a network address translation data store is not empty. In this example, a six byte source IP address and port can be used to attempt to retrieve information about a currently active network address translation connection. When the retrieval fails, the four byte source ip address can be used to attempt to find a network address translation initiator object which is stored as a collector object in the same complex data structure.
30 If an initiator is found, a connection is added to the data store under the six byte source

IP address and port retrieval path. At the same time, a connection object will be automatically added to the initiator object's member list because the initiator object is a collector object with a retrieval path that is the prefix of the connection object's retrieval path.

5 The list retrieval functions are used whenever the appliance needs to sequentially access a group of container objects. For example, a connection in a network address translation data store can be deleted when it has not been used for a length of time. In this case, an appliance could periodically step through part of the list of all connection objects in the data store to determine whether some connections
10 should be deleted.

 In another example, both Trie and List functions may be used when a request is received for information about all of the network address translation connection objects that were created using an initiator object identified by a four byte ip address. In this case, the initiator object is retrieved using the Trie retrieval function and a four byte
15 key. Then using the member list of the initiator object, the connection objects collected by the initiator are sequentially retrieved.

 FIGURE 4A illustrates a data object 110 that includes a pointer 112 to the actual location of the requested data. A parent list (not shown) identifies a parent object 114 that is associated with the data object 110. The parent list identifies each parent object that references the data object in each data structure (List and Trie). When the data
20 object 110 is subsequently deleted, the parent list is employed to automatically delete every parent object 114 in the List and Trie data structures that reference the deleted data object.

 FIGURE 4B shows a collector object 116, which is similar in some ways to the
25 data object 110 illustrated in FIGURE 4A, e.g., the collector object includes a data pointer 122 to the actual location of the requested data and a parent object 120 with a list identifying each reference to the collector object in both data structures. However, the collector object 116 further comprises a member object 118 with a list that identifies every other data object that is related to the collector object. When the data associated
30 with the collector object 116 is retrieved, the invention employs the member object 118

to also automatically retrieve the data associated with the other related data objects on the member list.

Additionally, each collector object automatically adds all of the objects (data and/or collector) to its member object from the sub-tree (node levels) in the Trie data structure below the reference to the collector object. Thus, when a collector object is referenced by a Trie leaf node, the collector object will have an empty list in its member object.

In one embodiment of the invention, an add function automatically adds an object to the list for the member object of a collector object when the added object is added beneath the existing level of the collector object in a sub-tree of the Trie data structure. Additionally, when another add function is employed to add a new collector object to the Trie data structure at a level above the levels of previously existing objects in a sub-tree of the Trie data structure, these objects are automatically added to the list in the member object associated with the new collector object.

FIGURE 5 shows an overview of one embodiment enabling a single collector object to reference other data objects and collector objects. A collector object 117 includes a member object (not shown) with a list that references all of the other data objects 110 and another collector object 119 disposed in a sub-tree below the reference to the collector object 117 in the Trie data structure. The other collector object 119 includes another member object (not shown) with another list that references the other data objects 110 disposed in another sub-tree below the reference to the collector object 119 in the Trie data structure. Thus, when the data associated with the collector object 117 is requested, other data associated with four other data objects 110 and another collector object 119 will be automatically retrieved.

FIGURE 6 illustrates an overview that provides greater detail regarding the referencing of data/collector objects in a data store 109 by an exemplary List data structure 107. Although the Functional Interface 102 and the Trie data structure 104 are represented in FIGURE 6, their interaction with the data store 108 and the List data structure 107 is discussed in greater detail above and below. A Port list 130 includes references to three data objects 110 which are associated with a particular port ("Y"),

but have different IP addresses. FIGURE 9B shows an exemplary Port list 184 that includes references to three separate data objects with the same port ("Y") and different IP addresses ("F", "G", and "H").

Returning to FIGURE 6, an IP address list 128 references two data and one collector objects that are associated with different Ports, but have the same IP address ("X"). FIGURE 9A illustrates an exemplary IP address list 182 that includes references to three separate data objects with the same IP address ("X") and different Ports ("U", "V", and "W").

FIGURE 7 illustrates an overview 132 that provides greater detail regarding the referencing of data/collector objects in a data store 111 by an exemplary Trie data structure 103. Although the Functional Interface 102 and the List data structure 106 are represented in FIGURE 7, their interaction with the data store 111 and the Trie data structure 103 is discussed in greater detail above and below. Each transition between each node 136 in the twelve levels of the Trie data structure 103 is associated with a four bit key segment having the decimal range of 0 to 15. In combination, these key segments can serve as the reference to the data object 110 in the data store 111. In this example, each of the two exemplary paths through the twelve Trie nodes correspond to the same IP address and different ports (Port W and Port V) that are associated with the same data object 110 in the data store 111.

In greater detail, FIGURE 10 shows two paths through eight levels of Trie nodes 186 that correspond to two different IP addresses and which reference the same data object 188. The sequential combination of each transition value between each of the eight Trie nodes 186 corresponds to a particular IP address. At each Trie node, four bits are employed to indicate incoming and outgoing transition pointers associated with any of 16 possible values for each transition between each node. For this example, the sequential combination of the eight transition values represent two 32 IP bit addresses, as follows: (4,8,0,15,6,0,3,12,9) and (4,8,0,15,6,8,15,0,7).

FIGURE 8 shows an overview 144 that provides greater detail regarding the referencing of exemplary data/collector objects in a data store 115 by an exemplary Trie data structure 117. Although the Functional Interface 102 and the List data structure

104 are represented in FIGURE 8, they are discussed in greater detail above and below. Each transition between each node 136 in the path through the twelve levels in the Trie data structure 117 is associated with a four bit key segment. In this case, the three complete paths through all twelve levels of the Trie nodes reference separate data objects associated with different ports and/or IP addresses, i.e., data object 138 is associated with a different port and IP address and data objects 139 and 140 are associated with the same IP address and different ports.

Additionally, a collector object 142 is referenced by the transition values associated with a path through the first four levels of the Trie data structure and which represent the first four values in a 32 bit IP address. In this case, the collector object 142 has automatically included a member object (not shown) that lists a reference to each data object (139 and 140) that is disposed in the sub-tree (node levels) below the reference to the collector object 142. Since the first four key segments of the collector object 142 and the data objects 139 and 140 are the same, these data objects were automatically included in the member object for the collector object 142.

FIGURE 11 illustrates a flow chart 146 that shows the actions for one embodiment of the invention for retrieving data in response to a request. Moving from a start block, the operational flow advances to a block 148, where the Functional interfaces receives a function request for data at an IP address and port. The operational flow pushes to a decision block 150, where a determination is made whether the function request is a List function. If so, the operational flow advances to a block 162 where the List data structure is employed to locate a data object at the provided port and IP address and retrieve the requested data. The operational flow moves to a block 164, where when the data object is determined to be a “collector” object, a member object associated with the collector object is employed to retrieve other related data identified by other data objects that are referenced by the member object. Next, the operational flow moves to the End block and resumes calling program modules.

Alternatively, when the determination at the decision block 150 is false, the flow moves to a block 152 where key segments are created for a Trie data structure based on the port and IP address associated with the requested data. In one embodiment of the

invention, each key segment is four bits. In this example, four key segments are created for a 16 bit port and eight key segments are generated for a 32 bit IP address.

The operational flow moves to a block 154 where each key segment is employed to sequentially transition along a path through all of the node levels in the Trie data structure. Advancing to a decision block 156, a determination is made whether a null pointer value was found at any level in the path through the nodes of the Trie data structure. If true, the operational flow moves to an End block and resumes calling program modules.

Alternatively, when the determination at the decision block 156 is false, the operational flow advances to a block 158 where the complete path through the Trie data structure is employed to reference a data object and retrieve the requested data. The operational flow moves to a block 160 where when the referenced data object is determined to also be a collector object, a member list included with the collector object is employed to retrieve related data associated with other data objects on the member list. The operational flow transitions to the end block and resumes calling program modules.

FIGURE 12 illustrates a flow chart 166 that shows the actions for one embodiment of the invention for deleting data in response to a request. Moving from a start block, the flow advances to a block 168, where a function request is received by the Functional Interface to delete data associated with a particular port and IP address. The operational flow transitions to a decision block 170, where the Functional Interface determines whether the data to be deleted is a collector object. If so, the operational flow advances to a block 182 where the list of member objects is removed from the collector object. The operational flow moves to a block 172 where a parent object associated with the data object to be deleted is identified. A list may be associated with the object to be deleted that identifies multiple parent objects. Also, since a parent object is automatically generated for each data object, it is automatically updated until the associated data object is deleted.

Alternatively, when the object to be deleted is determined to not be a collector object at the decision block 170, the operational flow moves directly to the block 172 where substantially the same actions discussed above are performed.

Next, the operation flow advances to a decision block 174 where a
5 determination is made whether the parent object is a trie node. If so, the operational flow moves to a block 184 where every trie node that references this parent object is deleted. The operational flow transitions to a decision block 178 where a determination is made whether more parent objects are included on the list associated with the data
10 object to be deleted. If affirmative, the operational flow loops back to the block 172 and performs substantially the same actions as discussed above. However, when the determination at the decision block 178 is negative, the operational flow moves to a block 180 where the data object to be deleted is deallocated. The operational flow transitions to an end block and resumes calling other program modules.

Alternatively, if the determination at the decision block 174 is negative, i.e., the
15 parent object is not a trie node, then the operational flow advances to a block 176 where the reference in the List data structure to the parent object is deleted. The operational flow moves to the decision block 178 where substantially the same actions discussed above are repeated.

One embodiment of the invention enables each data object in the data store to be
20 accessed by one or more keys in constant time when the Trie data structure is employed to resolve a request for a data object. Each data object in the data store may be accessed by zero or more sequential lists. For one embodiment of the invention, an iterator object is employed to support multiple concurrent list states in the List data structure.

It is envisioned that the data store can be a relational, object oriented or a
25 combination of the two, database. Also, the data space could be a data warehouse.

The above specification, examples and data provide a complete description of the manufacture and use of the composition of the invention. Since many embodiments of the invention can be made without departing from the spirit and scope of the invention, the invention resides in the claims hereinafter appended.

WE CLAIM:

1.

1. A method for employing a plurality of data structures to optimize the retrieval of at least one data object over a network, comprising:

(a) storing each data object in a data store, each data object in the data store being separately referenced in each of the plurality of data structures;

(b) in response to a request for one data object, automatically determining one of the plurality of data structures best suited to retrieve the one data object and employing the determined data structure to locate and retrieve the one data object from the data store;

(c) in response to a request for a plurality of related data objects, automatically determining another one of the plurality of data structures best suited to retrieve the plurality of related data objects and employing the determined other one of the plurality of data structures to locate and retrieve the plurality of related data objects from the data store; and

(d) in response to a request to delete at least one data object, automatically deleting each reference to each deleted data object in the plurality of data structures such that each subsequent request for each deleted data object will be unsuccessful.

2. The method of Claim 1, further comprising:

(a) associating a parent object with each data object, the parent object identifying each reference for the associated data object in the plurality of data structures; and

(b) when at least one data object is deleted, employing each parent object associated with each deleted data object to identify each reference for the deleted

data object in the plurality of data structures such that each reference to the deleted data object is deleted.

3. The method of Claim 1, wherein the plurality of related data objects have at least one related characteristic, including port, IP address and type.

4. The method of Claim 1, wherein the plurality of data structures include List, Hash and Trie.

5. The method of Claim 4, wherein the data object is a collector object that is associated with a member object that identifies one or more other data objects that are referenced in a sub-tree below a reference to the requested collector object in the Trie data structure, the member object being employed to reference and retrieve each other data object when the collector object is retrieved.

6. The method of Claim 5, further comprising automatically enabling the member object to identify a new data object that is added to the sub-tree below the reference to the collector object.

7. The method of Claim 4, wherein the one of the plurality of data structures is the Trie data structure, further comprising:

- (a) identifying a key in the request for the data object;
- (b) dividing the key into segments; and
- (c) employing each segment to search the Trie data structure and locate the requested data object.

8. The method of Claim 7, wherein the key represents an IP address.

9. The method of Claim 7, wherein the key represents a port.

10. The method of Claim 7, wherein each segment is represented by at least one bit.

11. The method of Claim 1, wherein the data store is a database having a type that includes relational, object-oriented or a combination of relational and object oriented.

12. The method of Claim 1, wherein the data object is a container.

13. The method of Claim 11, wherein the data store is a data warehouse.

14. In a client-server operating environment, a method for employing a plurality of data structures to optimize the retrieval of at least one data object over a network for any one of a plurality of applications running in the client-server operating environment, comprising the acts of:

(a) enabling a server to store each data object in a data store, each data object in the data store being separately referenced in each of the plurality of data structures;

(b) when one data object is requested by a client, enabling the server to automatically choose one of the plurality of data structures best suited to fulfill the request and retrieve the one data object for the client;

(c) when a plurality of related data objects are requested by the client, enabling the server to automatically choose another one of the plurality of data structures best suited to locate and retrieve the plurality of related data objects for the client; and

(d) in response to a request to delete at least one data object from the client, enabling the server to automatically delete each reference to each deleted data object in the plurality of data structures such that each subsequent request for each deleted data object will be unsuccessful.

15. In client-server operating environment of Claim 14, further comprising the actions of:

(e.) enabling the server to associate a parent object with each data object, the parent object identifying each reference for the associated data object in the plurality of data structures; and

(b) when at least one data object is deleted, enabling the server to employ each parent object associated with each deleted data object to identify each reference for the deleted data object in the plurality of data structures such that each reference to the deleted data object is deleted.

16. The client server operating environment of Claim 14, wherein the one of the plurality of data structures is a Trie data structure, further comprising:

(a) enabling the server to identify a key in the request for the data object;

(b) enabling the server to divide the key into segments; and

(c) enabling the server to employ each segment to search the trie data structure and locate the requested data object.

17. In a server array controller operating environment, a method for employing a plurality of data structures to optimize the retrieval of at least one data object over a network for any one of a plurality of applications running in the client-server operating environment, comprising the acts of:

(a) enabling a server array controller to store each data object in a data store, each data object in the data store being separately referenced in each of the plurality of data structures;

(b) when one data object is requested by a network device, enabling the server array controller to automatically choose one of the plurality of data structures

best suited to fulfill the request and retrieve the location of the one data object on at least one server;

(c) when a plurality of related data objects are requested by the network device, enabling the server array controller to automatically choose another one of the plurality of data structures best suited to retrieve the location of the plurality of related data objects on at least one server; and

(d) in response to a request to delete at least one data object by the network device, enabling the server array controller to automatically delete each reference to each deleted data object in the plurality of data structures such that each subsequent request for each deleted data object will be unsuccessful.

18. In the server array controller operating environment of Claim 17, further comprising the actions of:

(a) enabling the server array controller to associate a parent object with each data object, the parent object identifying each reference for the associated data object in the plurality of data structures; and

(b) when at least one data object is deleted, enabling the server array controller to employ each parent object associated with each deleted data object to identify each reference for the deleted data object in the plurality of data structures such that each reference to the deleted data object is deleted.

19. The server array controller operating environment of Claim 17, wherein the plurality of data structures include Trie, Hash and List.

20. The server array controller operating environment of Claim 17, wherein the network device includes, router, client, cache, firewall and another server array controller.

21. A computer readable medium readable by a computing system and encoding a computer program of instructions for executing a computer process for

employing a plurality of data structures to optimize the retrieval of data over a network for an application on the computing system, comprising the actions of :

(a) storing each data object in a data store, each data object in the data store being separately referenced in each of the plurality of data structures;

(b) when one data object is requested, automatically determining one of the plurality of data structures best suited to retrieve the one data object and employing the determined data structure to locate and retrieve the one data object from the data store;

(c) when a plurality of data objects are requested, automatically determining another one of the plurality of data structures best suited to retrieve the plurality of related data objects and employing the determined other one of the plurality of data structures to locate and retrieve the plurality of related data objects from the data store; and

(d) in response to a request to delete at least one data object, automatically deleting each reference to each deleted data object in the plurality of data structures such that each subsequent request for each deleted data object will be unsuccessful.


22. A modulated data signal having computer executable instructions embodied thereon for employing a plurality of data structures to optimize the retrieval of data objects over the Internet, comprising :

(a) a transmitter for sending a request for one data object from a data store, each data object in the data store being separately referenced in each of the plurality of data structures;

(b) a receiver for receiving the request and enabling the automatic determination of one of the plurality of data structures best suited to locate and retrieve the one data object;

Abstract

The invention provides for employing a complex data structure to optimize the retrieval of data from a data store over a network. The complex data structure includes two separate sub-data structures (Trie and List) that separately reference the same data objects in a data store. The complex data structure employs a functional interface to determine which data structure matches a particular function request for data. A Trie sub-data structure is used to fulfill a single data object request. The List sub-data structure is employed with function requests related to several data objects. Each data object is associated with a parent object that includes a list of every reference to the data object in both the Trie and List sub-data structures. When a data object is subsequently deleted, the parent object list is employed to automatically delete every reference to the deleted data object in both the Trie and List sub-data structures. A collector object is a type of data object that is associated with a member object and which includes a list of other related data/collector objects. When data associated with the collector object is requested, other data associated with the other data/collector objects on the member object list are automatically retrieved. The complex data structure can also use a Hash sub-data structure to retrieve single data objects.

CERTIFICATE UNDER 37 CFR 1.10:	
"Express Mail" mailing label number:	EL626731366US
Date of Deposit:	June 19, 2000
I hereby certify that: this paper or fee is being deposited with the U.S. Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to BOX PATENT APPLICATION, Assistant Commissioner for Patents, Washington, D.C. 20231.	
By:	
Name:	Janice Knierp




FIG. 1A

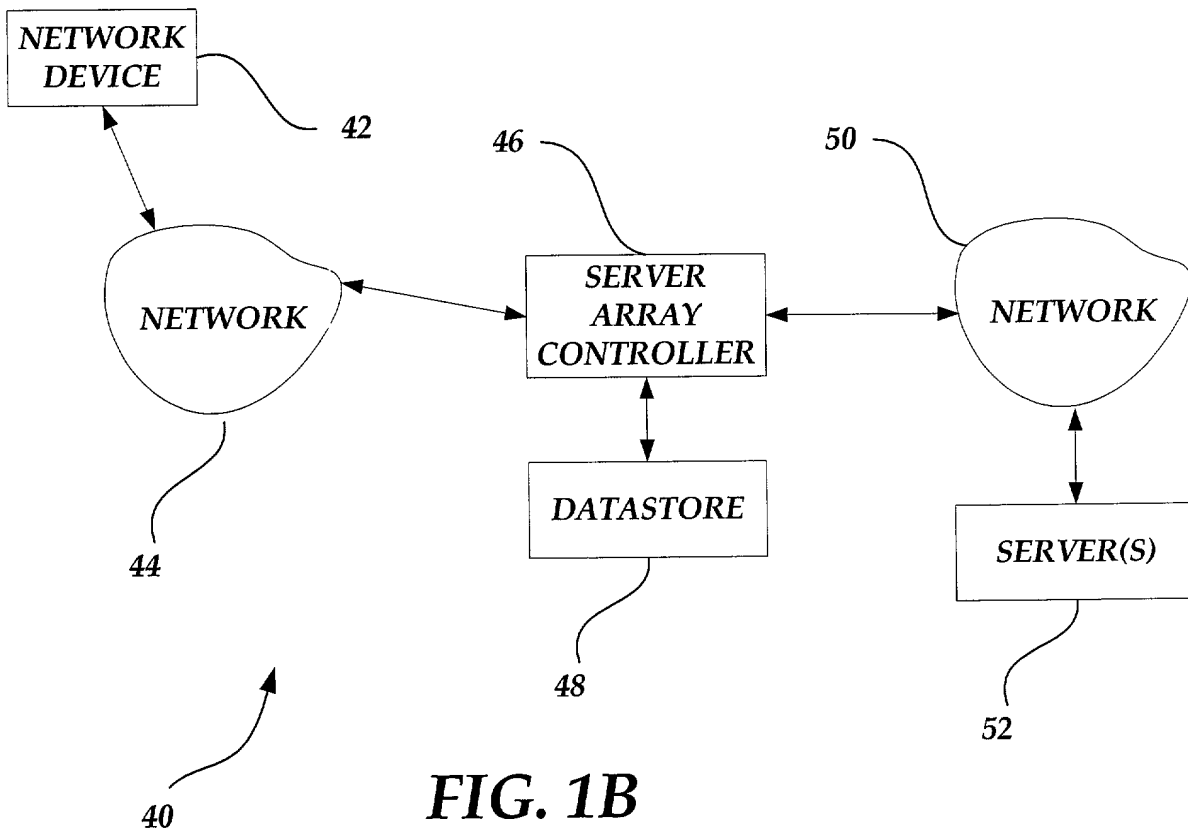
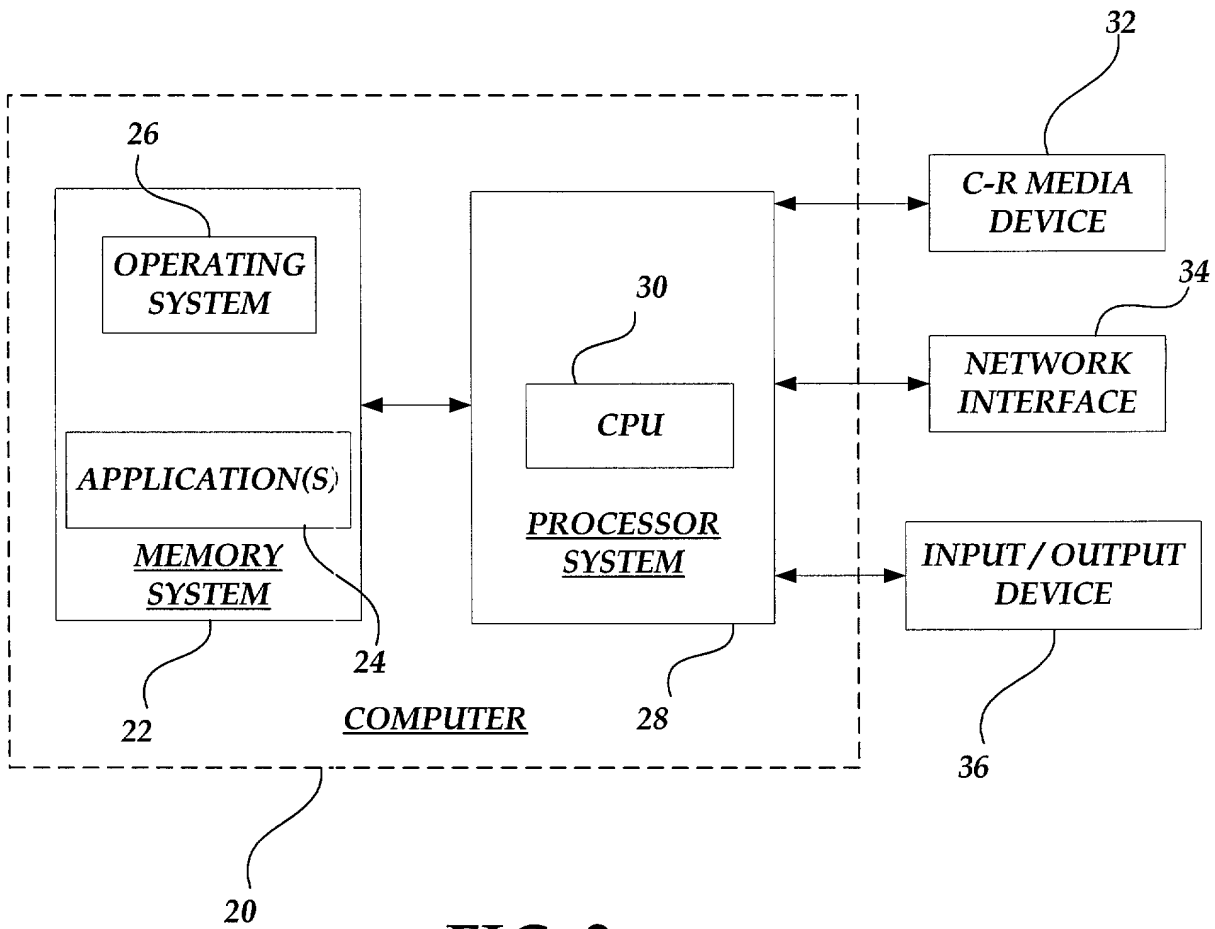



FIG. 1B

**FIG. 2**



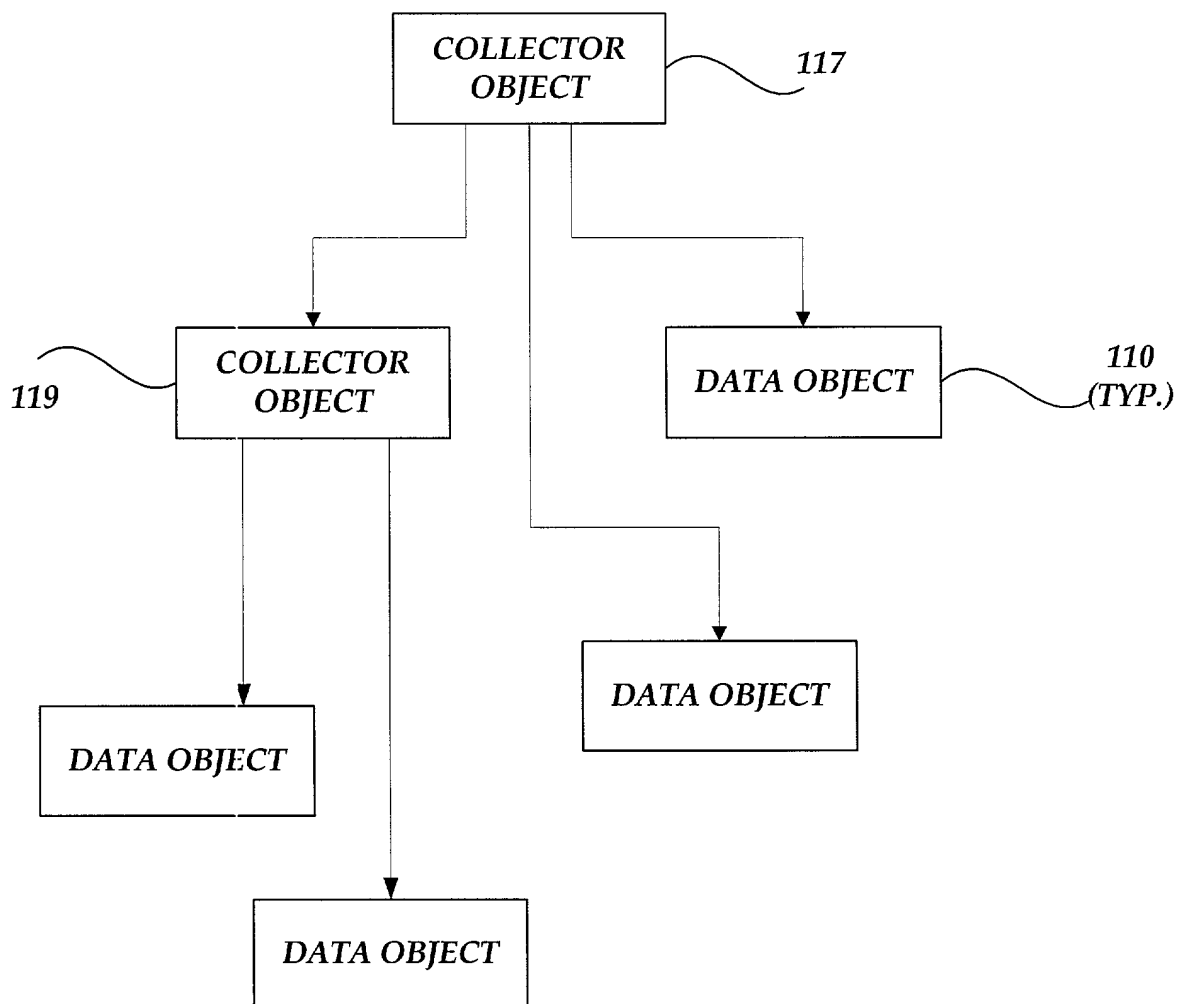
100



FIG. 4A



FIG. 4B

**FIG. 5**

124

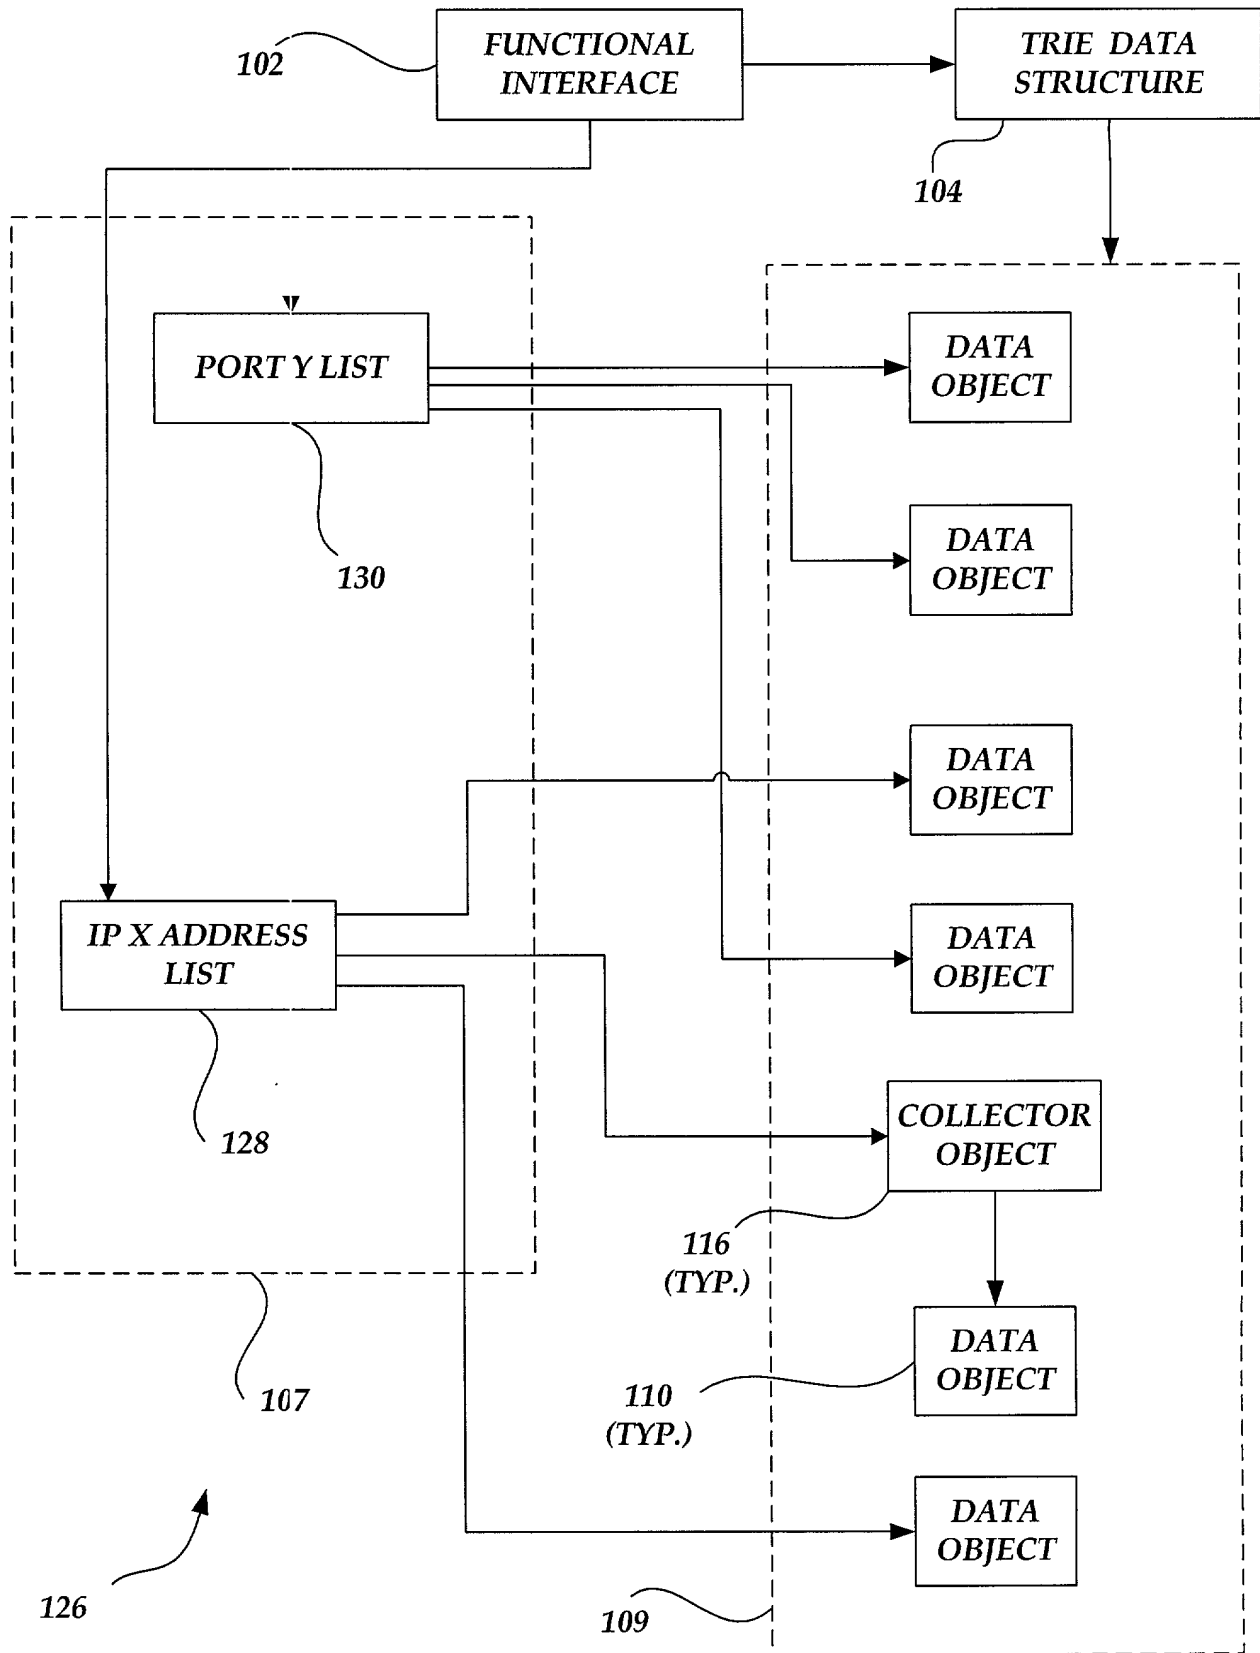


FIG. 6

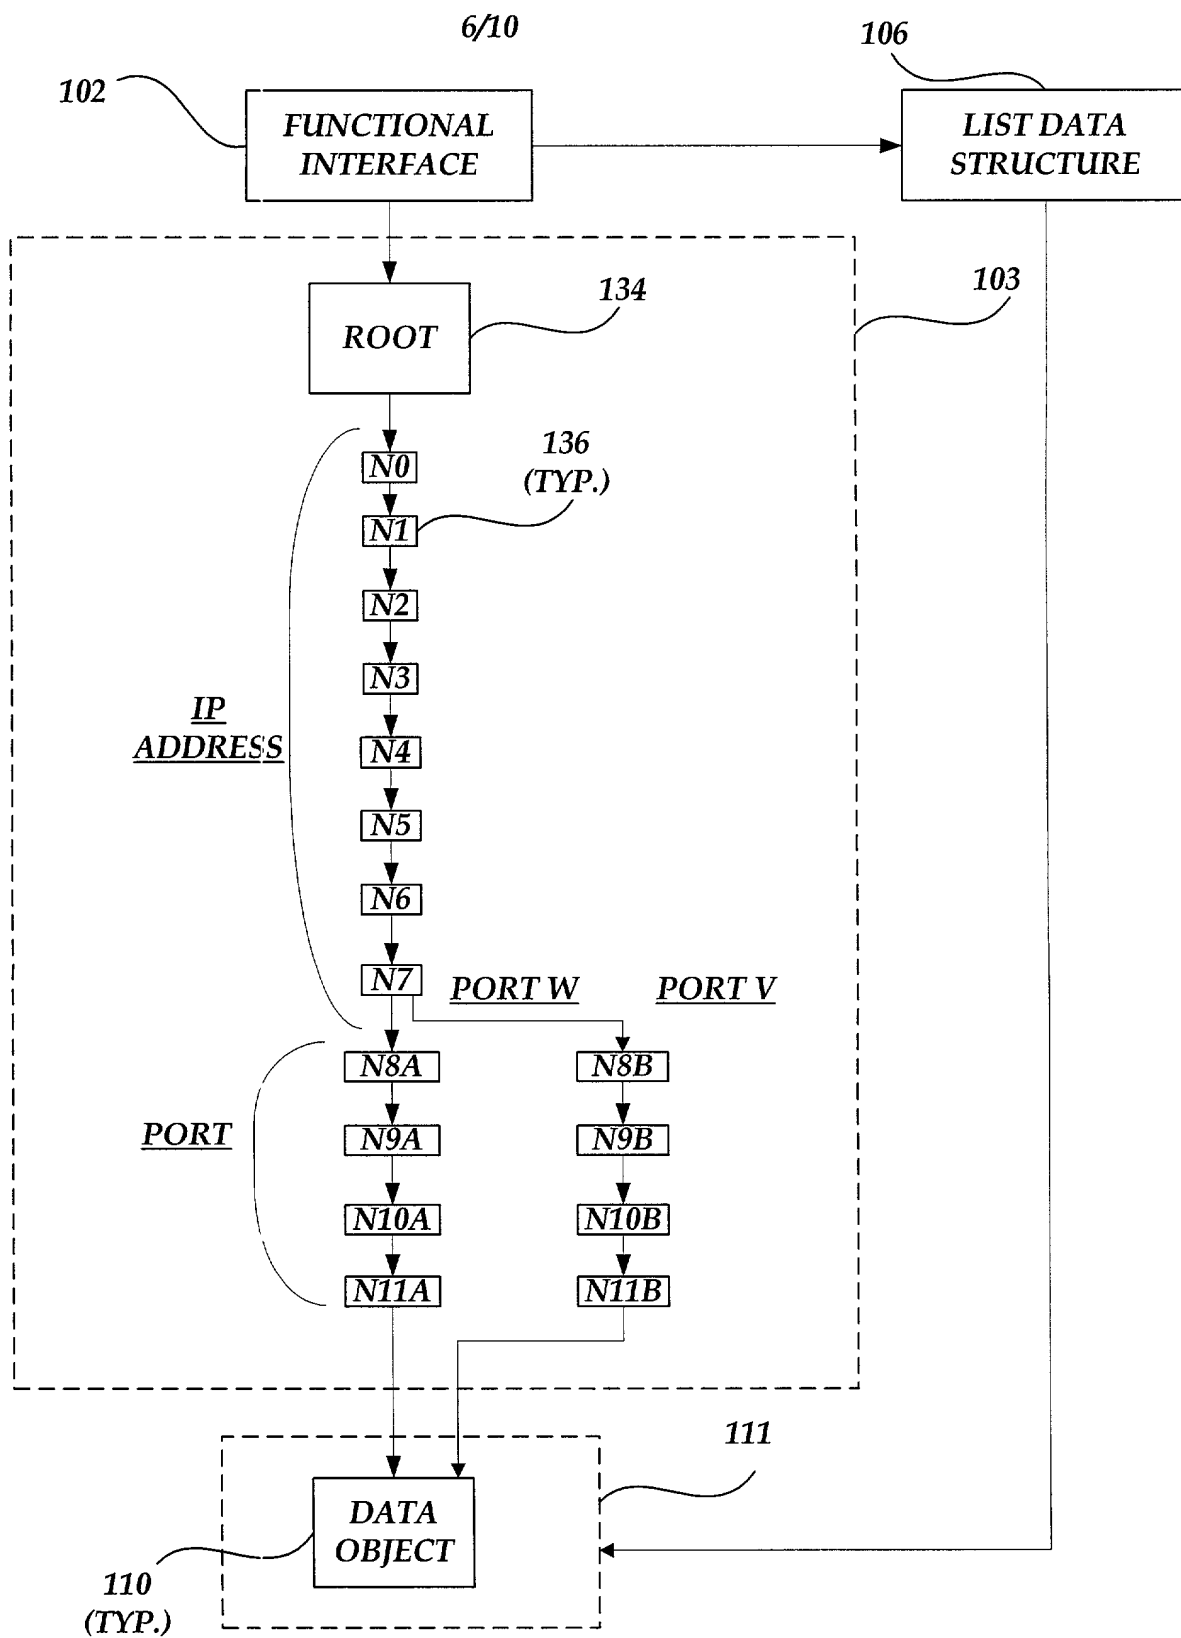


FIG. 7

132




FIG. 8





FIG. 9A

FIG. 9B


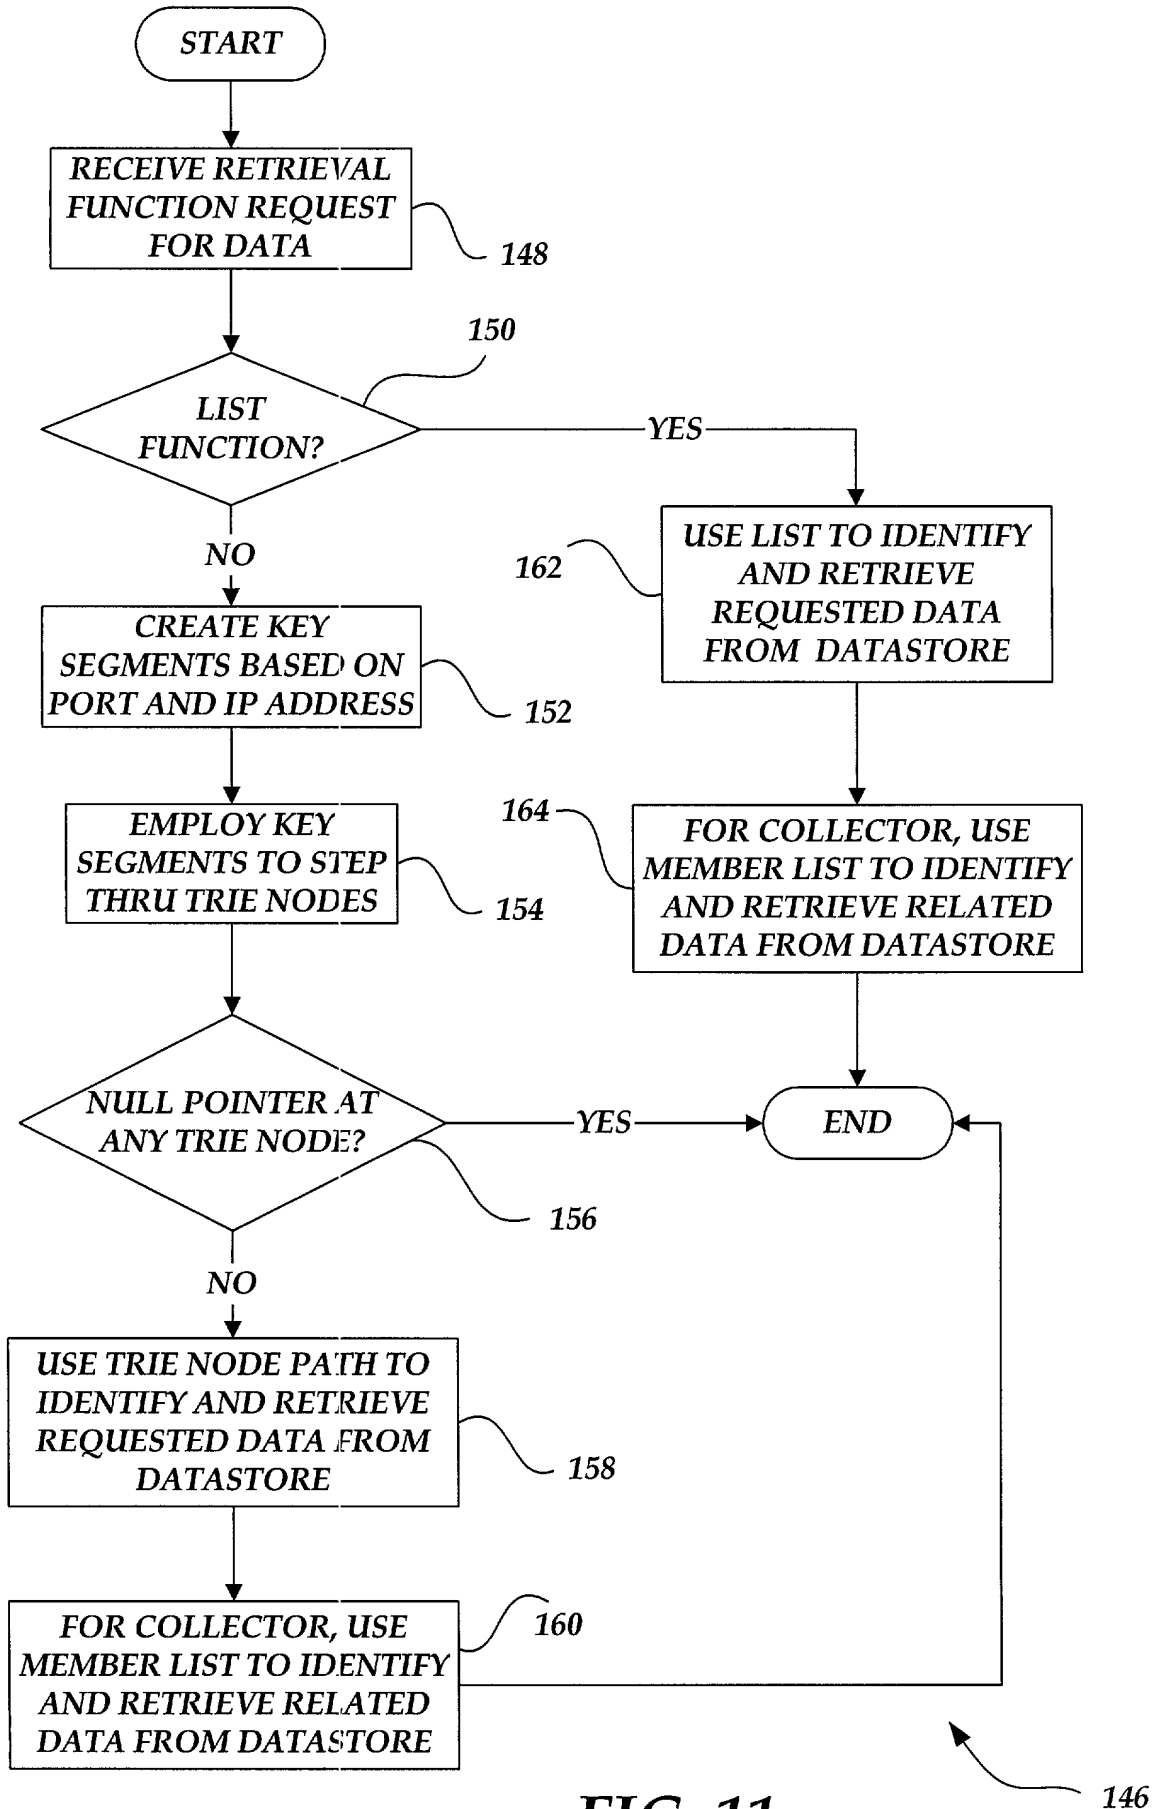


FIG. 10



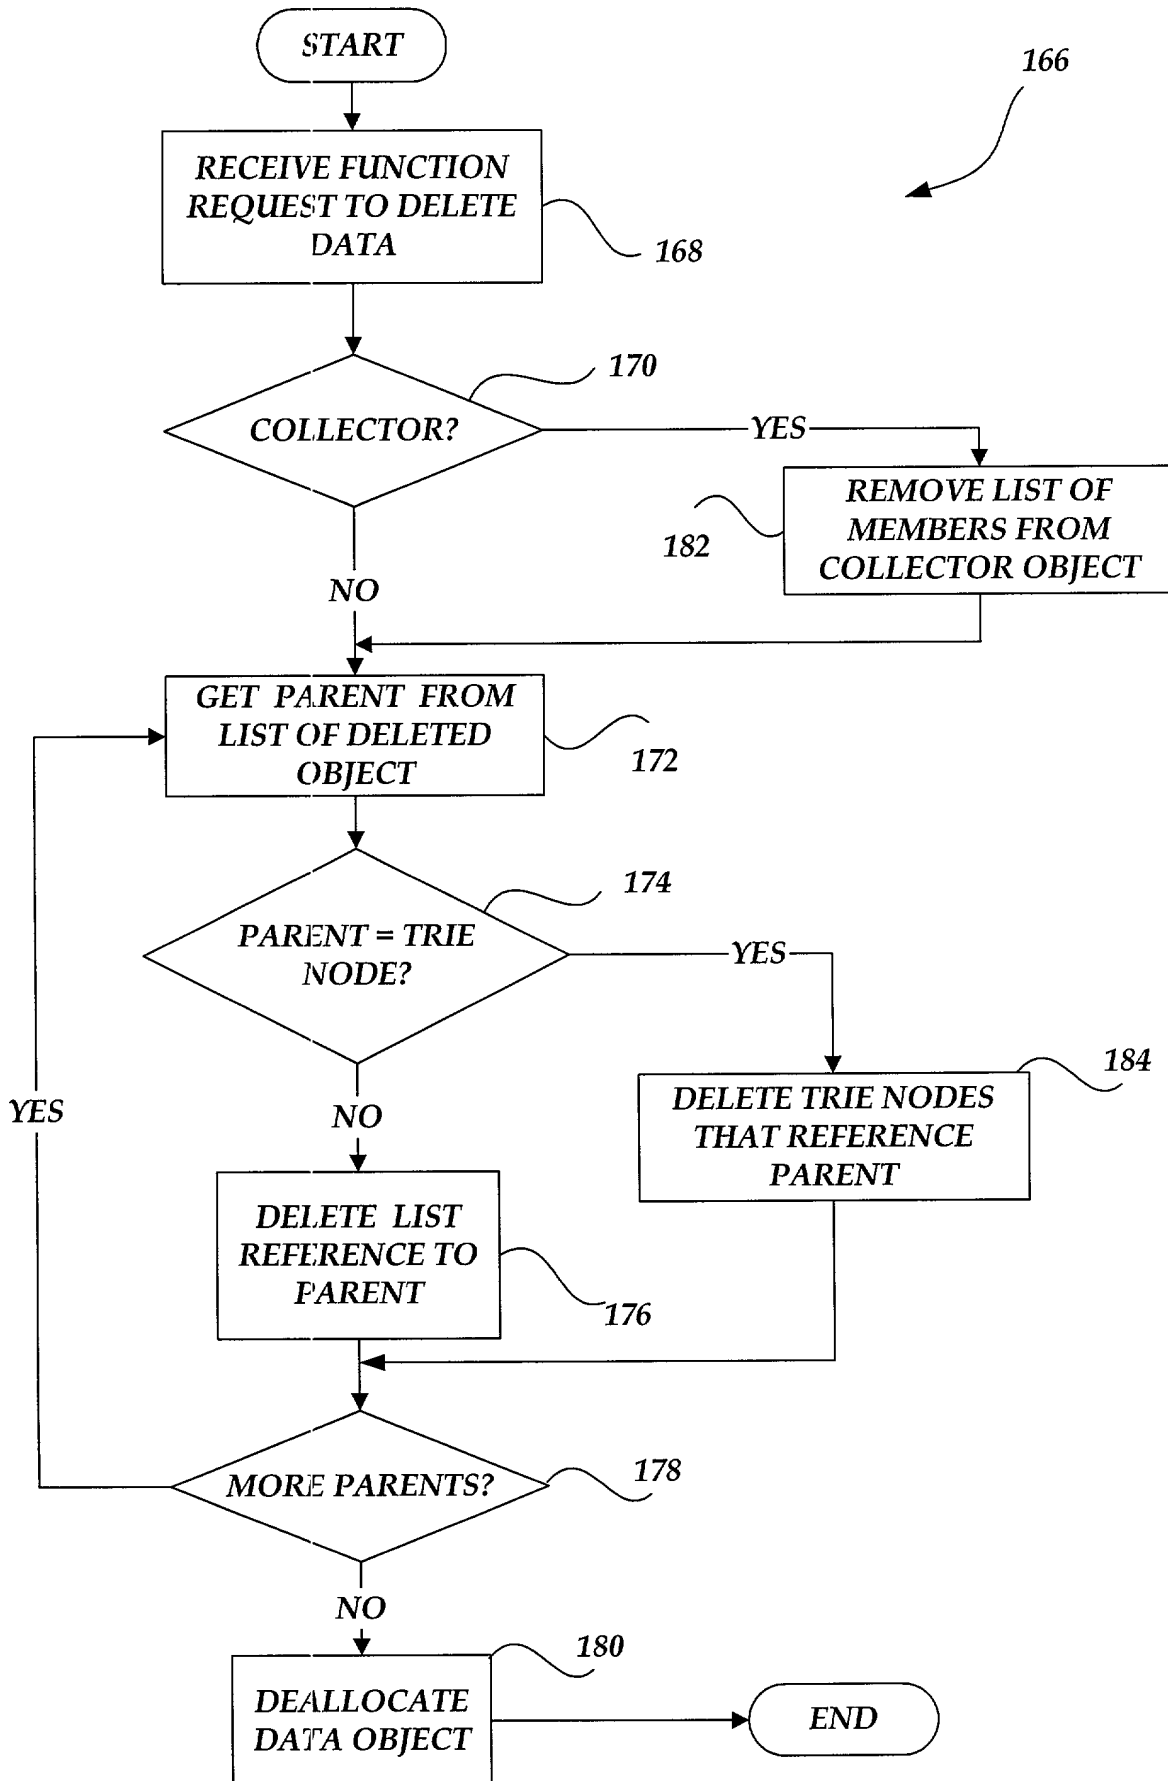


FIG. 12

MERCHANT & GOULD P.C.

United States Patent Application

COMBINED DECLARATION AND POWER OF ATTORNEY

As a below named inventor I hereby declare that: my residence, post office address and citizenship are as stated below next to my name; that

I verily believe I am the original, first and sole inventor (if only one name is listed below) or a joint inventor (if plural inventors are named below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: METHOD AND SYSTEM FOR NETWORK LOAD BALANCING WITH A COMPOUND DATA STRUCTURE

The specification of which

- a. ☒ is attached hereto
b. ☐ was filed on _____ as application serial no. _____ and was amended on _____ (if applicable) (in the case of a PCT-filed application) described and claimed in international no. _____ filed _____ and as amended on _____ (if any), which I have reviewed and for which I solicit a United States patent.

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, § 1.56 (attached hereto).

I hereby claim foreign priority benefits under Title 35, United States Code, § 119/365 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on the basis of which priority is claimed:

- a. ☐ no such applications have been filed.
b. ☒ such applications have been filed as follows:

FOREIGN APPLICATION(S), IF ANY, CLAIMING PRIORITY UNDER 35 USC § 119			
COUNTRY	APPLICATION NUMBER	DATE OF FILING (day, month, year)	DATE OF ISSUE (day, month, year)
ALL FOREIGN APPLICATION(S), IF ANY, FILED BEFORE THE PRIORITY APPLICATION(S)			
COUNTRY	APPLICATION NUMBER	DATE OF FILING (day, month, year)	DATE OF ISSUE (day, month, year)

I hereby claim the benefit under Title 35, United States Code, § 120/365 of any United States and PCT international application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, § 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, § 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application.

U.S. APPLICATION NUMBER	DATE OF FILING (day, month, year)	STATUS (patented, pending, abandoned)

I hereby claim the benefit under Title 35, United States Code § 119(e) of any United States provisional application(s) listed below:

U.S. PROVISIONAL APPLICATION NUMBER	DATE OF FILING (Day, Month, Year)
60/140,272	June 18, 2000

I hereby appoint the following attorney(s) and/or patent agent(s) to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith:

Albrecht, John W.	Reg. No. 40,481	Lacy, Paul E.	Reg. No. 38,946
Anderson, Gregg I.	Reg. No. 28,828	Larson, James A.	Reg. No. 40,443
Batzli, Brian H.	Reg. No. 32,960	Liepa, Mara E.	Reg. No. 40,066
Beard, John L.	Reg. No. 27,612	Lindquist, Timothy A.	Reg. No. 40,701
Berns, John M.	Reg. No. 43,496	Lycke, Lawrence E.	Reg. No. 38,540
Black, Bruce E.	Reg. No. 41,622	McAuley, Steven A.	Reg. No. 46,084
Branch, John W.	Reg. No. 41,633	McDonald, Daniel W.	Reg. No. 32,044
Bremer, Dennis C.	Reg. No. 40,528	McIntyre, Jr., William F.	Reg. No. P-44,921
Bruess, Steven C.	Reg. No. 34,130	Mueller, Douglas P.	Reg. No. 30,300
Byrne, Linda M.	Reg. No. 32,404	Pauly, Daniel M.	Reg. No. 40,123
Carlson, Alan G.	Reg. No. 25,959	Phillips, John B.	Reg. No. 37,206
Caspers, Philip P.	Reg. No. 33,227	Plunkett, Theodore	Reg. No. 37,209
Chiapetta, James R.	Reg. No. 39,634	Prendergast, Paul	Reg. No. 46,068
Clifford, John A.	Reg. No. 30,247	Pytel, Melissa J.	Reg. No. 41,512
Cochran, William W.	Reg. No. 26,652	Qualey, Terry	Reg. No. 25,148
Daignault, Ronald A.	Reg. No. 25,968	Reich, John C.	Reg. No. 37,703
Daley, Dennis R.	Reg. No. 34,994	Reiland, Earl D.	Reg. No. 25,767
DalGLISH, Leslie E.	Reg. No. 40,579	Schmaltz, David G.	Reg. No. 39,828
Daulton, Julie R.	Reg. No. 36,414	Schuman, Mark D.	Reg. No. 31,197
DeVries Smith, Katherine M.	Reg. No. 42,157	Schumann, Michael D.	Reg. No. 30,422
DiPietro, Mark J.	Reg. No. 28,707	Scull, Timothy B.	Reg. No. 42,137
Edell, Robert T.	Reg. No. 20,187	Sebald, Gregory A.	Reg. No. 33,280
Epp, Ryan, Sandra	Reg. No. 39,667	Skoog, Mark T.	Reg. No. 40,178
Glance, Robert J.	Reg. No. 40,620	Spellman, Steven J.	Reg. No. 45,124
Goggin, Matthew J.	Reg. No. 44,125	Stoll-DeBell, Kirstin L.	Reg. No. 43,164
Golla, Charles E.	Reg. No. 26,896	Sumner, John P.	Reg. No. 29,114
Gorman, Alan G.	Reg. No. 38,472	Swenson, Erik G.	Reg. No. 45,147
Gould, John D.	Reg. No. 18,223	Tellekson, David K.	Reg. No. 32,314
Gregson, Richard	Reg. No. 41,804	Trembath, Jon R.	Reg. No. 38,344
Greshens, John J.	Reg. No. 33,112	Underhill, Albert L.	Reg. No. 27,403
Hamre, Curtis B.	Reg. No. 29,165	Vandenburgh, J. Derek	Reg. No. 32,179
Hillson, Randall A.	Reg. No. 31,838	Wahl, John R.	Reg. No. 33,044
Holzer, Jr., Richard J.	Reg. No. 42,668	Weaver, Karrie G.	Reg. No. 43,245
Johnston, Scott W.	Reg. No. 39,721	Welter, Paul A.	Reg. No. 20,890
Kachewitch, Natalie D.	Reg. No. 34,196	Whipps, Brian	Reg. No. 43,261
Karjeker, Shaukat	Reg. No. 34,049	Wickhem, J. Scot	Reg. No. 41,376
Kastelic, Joseph M.	Reg. No. 37,160	Williams, Douglas J.	Reg. No. 27,054
Kettelberger, Denise	Reg. No. 33,924	Witt, Jonelle	Reg. No. 41,980
Keys, Jeramie J.	Reg. No. 42,724	Wu, Tong	Reg. No. 43,361
Knearl, Homer L.	Reg. No. 21,197	Xu, Min S.	Reg. No. 39,536
Kowalchyk, Alan W.	Reg. No. 31,535	Zeuli, Anthony R.	Reg. No. 45,255
Kowalchyk, Katherine M.	Reg. No. 36,848		

I hereby authorize them to act and rely on instructions from and communicate directly with the person/assignee/attorney/firm/ organization who/which first sends/sent this case to them and by whom/which I hereby declare that I have consented after full disclosure to be represented unless/until I instruct Merchant & Gould P.C. to the contrary.

Please direct all correspondence in this case to Merchant & Gould P.C. at the address indicated below:

Merchant & Gould P.C.
P.O. Box 2903
Minneapolis, MN 55402-0903



I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

201	Full Name Of Inventor	Family Name Tennican	First Given Name Scott	Second Given Name P.
	Residence & Citizenship	City Seattle	State or Foreign Country Washington	Country of Citizenship United States
	Post Office Address	Post Office Address 7539 40 th Avenue NE	City Seattle	State & Zip Code/Country Washington 98115 / USA
Signature of Inventor 201:				Date:

202	Full Name Of Inventor	Family Name Skene	First Given Name Bryan	Second Given Name D.
	Residence & Citizenship	City Seattle	State or Foreign Country Washington	Country of Citizenship United States
	Post Office Address	Post Office Address 3026 NE 103 rd Street	City Seattle	State & Zip Code/Country Washington 98125 / USA
Signature of Inventor 202:				Date:

§ 1.56 Duty to disclose information material to patentability.

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclose information exists with respect to each pending claim until the claim is canceled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is canceled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclose all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§ 1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:

(1) prior art cited in search reports of a foreign patent office in a counterpart application, and

(2) the closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.

(b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made of record in the application, and

(1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim;

(2) It refutes, or is inconsistent with, a position the applicant takes in:

(i) Opposing an argument of unpatentability relied on by the Office, or

(ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

(1) Each inventor named in the application:

(2) Each attorney or agent who prepares or prosecutes the application; and

(3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.

(d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent, or inventor.